

Claims

1. A method for handling dynamic state information used for handling data packets, which arrive at a network element node of a network element cluster, said network element cluster having at least two nodes and each node handling separate sets of data packets, said method comprising the step of:
 - maintaining (206) in a node a first, node-specific data structure (557, 558, 559) comprising entries representing state information (520) needed for handling sets of data packets handled in said node,
characterized in that said method further comprises the step of:
 - maintaining (208) in said node in addition to said node-specific data structure a second, common data structure (554, 555, 556) comprising at least entries representing state information (520) needed for handling sets of data packets handled in one other node of said network element cluster, the contents of said common data structure effectively differing from the contents of said node-specific data structure.
2. A method according to claim 1, **characterized** in that it further comprises the steps of:
 - allocating (200) to each node belonging to said network element cluster certain node-specific distribution identifiers, each node having separate node-specific distribution identifiers allocated to it,
 - handling at least a plurality of data packets so that a data packet is handled (204) in that node of said network element cluster, to which node a distribution identifier calculated (202) using certain field(s) of said data packet is allocated, and
 - maintaining (212) in a plurality of entries of said node-specific and common data structures distribution information (510) relating to the distribution identifier, which corresponds to the set of data packets related to the respective entry.
3. A method according to claim 2, **characterized** in that it further comprises the steps of:
 - reallocating (605, 606, 607) said distribution identifiers to the nodes of said network element cluster,
 - if said reallocation results in a new distribution identifier being allocated to a node, said new distribution identifier being a distribution identifier not allocated to said node at the time of the reallocation, identifying (612) in the common data structure of said node the entries corresponding to said new distribution identifier, and adding (613) said entries to the node-specific data structure of said node, and

- if said reallocation results in an old distribution identifier not being allocated to a node anymore, said old distribution identifier being a distribution identifier allocated to said node at the time of the reallocation, identifying (615) in the node-specific data structure of said node the entries corresponding to said old distribution identifier, and clearing (616) said entries from the node-specific data structure of said node.

4. A method according to claim 2, **characterized** in that it further comprises the steps of:

- adding (400) a new entry to said node-specific data structure in a first node,
- communicating (402) said new entry at least to a second node of the network element cluster, and
- adding (403) an entry corresponding to said new entry to the common data structure of said second node.

5. A method according to claim 4, **characterized** in that it further comprises the step of:

- adding (401) an entry corresponding to said new entry to the common data structure of said first node.

6. A method according to claim 1, **characterized** in further maintaining (210) in said common data structure of said node entries representing state information needed for handling sets of data packets handled in said node.

7. A method according to claim 1, **characterized** in that said state information comprises the source address field (521a) and/or the destination address field (521b) of an Internet Protocol header, and/or port header fields (522a, 522b) of a Transmission Control Protocol header and/or port header fields (522a, 522b) of a User Datagram Protocol header, and/or the identifier header field of an Internet Control Message Protocol header, and/or a Message Identifier field (524) of an Internet Security Association and Key Management Protocol header, and/or an Initiator Cookie field (525) of an Internet Security Association and Key Management Protocol header, and/or the Security Parameter Index field (523) of a security header relating to the IPSec protocol suite, and/or a Session ID field (526) relating to the Secure Sockets Layer protocol, and/or an HTTP Cookie field (527) relating to the HyperText Transfer Protocol.

8. A method according to claim 1, **characterized** in that said state information comprises information (528) identifying an authenticated entity.

9. A method according to claim 1, **characterized** in that said state information comprises information (523) identifying a secured tunnel, within which data packets of the corresponding set are tunneled.
10. A method according to claim 2, **characterized** in that said distribution identifier is a hash value (512) and a hash function is used for calculating a hash value using certain field(s) of a data packet.
11. A method according to claim 2, **characterized** in that said distribution information is said distribution identifier (511).
12. A method according to claim 2, **characterized** in that said distribution information is information needed for calculating said distribution identifier for the corresponding data packet.
13. A method according to claim 2, **characterized** in that said certain field(s) for calculating a distribution identifier comprise the source address field (521a) and/or the destination address field (522b) of an Internet Protocol header, and/or port header fields (522a, 522b) of a Transmission Control Protocol header and/or port header fields (522a, 522b) of a User Datagram Protocol header, and/or the identifier header field of an Internet Control Message Protocol header, and/or a Message Identifier field (524) of an Internet Security Association and Key Management Protocol header, and/or an Initiator Cookie field (525) of an Internet Security Association and Key Management Protocol header, and/or the Security Parameter Index field (523) of a security header relating to the IPSec protocol suite, and/or a Session ID field (526) relating to the Secure Sockets Layer protocol, and/or an HTTP Cookie field (527) relating to the HyperText Transfer Protocol.
14. A network element node (700) of a network element cluster having at least two nodes, said node (700) comprising
 - first storage means (704), and
 - means (702) for maintaining in said first storage means (704) a first, node-specific data structure (551, 552, 553) comprising entries representing state information (520) needed for handling sets of data packets handled in said node,**characterized** in that said node further comprises:
 - second storage means (708), and
 - means (706) for maintaining in said second storage means (708) a second, common data structure (554, 555, 556) comprising at least entries representing state information needed for handling sets of data packets handled in one other node of

said network element cluster, the contents of said common data structure effectively differing from the contents of said node-specific data structure.

15. A network element node (700) according to claim 14, **characterized** in that:

- said means (702) for maintaining the node-specific data structure are adapted to add a new entry to said node-specific data structure in said first storage means (704), and to communicate said new entry to said means (706) for maintaining common data structure,

- said means (706) for maintaining the common data structure are adapted to communicate said new entry at least to one other node of the network element cluster, and in that

- said means (706) for maintaining the common data structure are further adapted to receive an entry from at least one other node of the network element cluster and to add an entry corresponding to said received entry to said common data structure in said second storage means (708).

16. A network element node (700) according to claim 15, **characterized** in that:

- said means (706) for maintaining the common data structure are further adapted to add a new entry received from said means (702) for maintaining the node-specific data structure to said common data structure in said second storage means (708).

17. A network element node (700) according to claim 14, **characterized** in that it further comprises:

- means (710) for receiving distribution identifiers, which are currently allocated to said node, said distribution identifiers being used for handling at least a plurality of data packets so that a data packet is handled in that node of said network element cluster, to which node a distribution identifier calculated using certain field(s) of said data packet is allocated, and

- third storage means (712) for storing said distribution identifiers, and in that

- said means (702, 706) for maintaining the node-specific and common data structures are adapted to maintain in a plurality of entries of said node-specific and common data structures in said first and second storage means (704, 708) distribution information relating to the distribution identifier, which corresponds to the set of data packets related to the respective entry.

18. A network element node according to claim 17, **characterized** in that:

- said means (710) for receiving distribution identifiers are adapted to receive reallocated distribution identifiers, and

- said means (706) for maintaining the common data structure are adapted to detect a new distribution identifier being allocated to said node due to the reallocation, said new distribution identifier being a distribution identifier not allocated to said node at the time of receiving reallocated distribution identifiers, and to identify in the common data structure the entries corresponding to said new distribution identifier, and to communicate said entries to said means (702) for maintaining the node-specific data structure for said entries to be added to the node-specific data structure, and
 - said means (702) for maintaining the node-specific data structure are adapted to detect an old distribution identifier not being anymore allocated to said node due to the reallocation, said old distribution identifier being a distribution identifier allocated to said node at the time of the reallocation, and to identify in the node-specific data structure the entries corresponding to said old distribution identifier, and to clear said entries from the node-specific data structure.
19. A network element node (700) according to claim 14, **characterized** in that said first storage means (704) is a portion of kernel space memory.
20. A network element node (700) according to claim 14, **characterized** in that said second storage means (708) is a portion of user space memory.
21. A network element node (700) according to claim 14, **characterized** in that said first storage means (704) is a portion of content addressable memory.
22. A network element node (700) according to claim 14, **characterized** in that said first storage means (704) part of a cryptographic card.
23. A network element cluster (800) having at least two network element nodes (700), at least one of said nodes (700) comprising
- first storage means (704), and
 - means (702) for maintaining in said first storage means (704) a first, node-specific data structure (551, 552, 553) comprising entries representing state information needed for handling sets of data packets handled in said node, **characterized** in that said at least one node further comprises:
 - second storage means (708), and
 - means (706) for maintaining in said second storage means (708) a second, common data structure (554, 555, 556) comprising at least entries representing state information needed for handling sets of data packets handled in one other node of

said network element cluster, the contents of said common data structure effectively differing from the contents of said node-specific data structure.

24. A network element cluster (800) according to claim 23, **characterized** in that it further comprises:

- 5 - means (802) for allocating to each node belonging to said network element cluster certain node-specific distribution identifiers, each node having separate node-specific distribution identifiers allocated to it, said distribution identifiers being used for handling at least a plurality of data packets so that a data packet is handled in that node of said network element cluster, to which node a distribution identifier
- 10 calculated using certain field(s) of said data packet is allocated, and in that said at least one node further comprises:
 - means (710) for receiving distribution identifiers, which are currently allocated to said node, and
 - third storage means (712) for storing said distribution identifiers, and in that
 - 15 - said means (702, 706) for maintaining the node-specific and common data structures are adapted to maintain in a plurality of entries of said node-specific and common data structures in said first and second storage means (704, 708) distribution information relating to the distribution identifier, which corresponds to the set of data packets related to the respective entry.
- 20 25. A network element cluster () according to claim 24, **characterized** in that:
 - said means (802) for allocating distribution identifiers are adapted to reallocate distribution identifiers to the nodes of said network element cluster, and in that in said at least one node
 - said means (710) for receiving distribution identifiers are adapted to receive
 - 25 reallocated distribution identifiers, and
 - said means (706) for maintaining the common data structure are adapted to detect a new distribution identifier being allocated to said node due to the reallocation, said new distribution identifier being a distribution identifier not allocated to said node at the time of receiving reallocated distribution identifiers, and to identify in
 - 30 the common data structure the entries corresponding to said new distribution identifier, and to communicate said entries to said means (702) for maintaining the node-specific data structure for said entries to be added to the node-specific data structure, and
 - said means (702) for maintaining the node-specific data structure are adapted to
 - 35 detect an old distribution identifier not being anymore allocated to said node due to the reallocation, said old distribution identifier being a distribution identifier

allocated to said node at the time of the reallocation, and to identify in the node-specific data structure the entries corresponding to said old distribution identifier, and to clear said entries from the node-specific data structure.

26. A computer program comprising program code for performing all the steps of
5 Claim 1 when said program is run on a computer.

27. A computer program product comprising program code means stored on a computer readable medium for performing the method of Claim 1 when said program product is run on a computer.